UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/479,146	01/07/2000	STEPHEN FULD	99-051-TAP	2688

7590 12/01/2004

TIMOTHY R SCHULTE
STORAGE TECHNOLOGY CORPORATION
ONE STORAGETEK DRIVE MS 4309
LOUISVILLE, CO 800284309

EXAMINER	
MASKULINSKI, MICHAEL C	
ART UNIT	**PAPER NUMBER**
2113	25

DATE MAILED: 12/01/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

3

PTO-90C (Rev. 10/03)

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

DEC 01 2004

Technology Center 2100

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Paper No. 25

Application Number: 09/479,146
Filing Date: January 07, 2000
Appellant(s): FULD, STEPHEN

Timothy R. Schulte
For Appellant

SUPPLEMENTAL EXAMINER'S ANSWER

This is in response to the reply brief filed September 29, 2003, paper no. 21.

(1) Real Party in Interest

A statement identifying the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

A statement identifying the related appeals and interferences, which will directly affect or be directly affected by or have a bearing on the decision in the pending appeal is contained in the brief.

(3) Status of Claims

The statement of the status of the claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

(5) Summary of Invention

The summary of invention contained in the brief is correct.

(6) Issues

The appellant's statement of the issues in the brief is correct.

(7) Grouping of Claims

Appellant's brief includes a statement that claims 12-14 and 16-18 stand or fall together.

(8) Claims Appealed

The copy of the appealed claims contained in the Appendix to the brief is correct.

(9) Prior Art of Record

6,018,778 Stolowitz 1-2000

White, Ron, How Computers Work, Millennium Edition, 1999, Que Corporation, pp. 176-177

(10) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claims 12-14 and 16-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Stolowitz, U.S. Patent 6,018,778, and further in view of White, How Computers Work.

Referring to the limitation "the storage elements are magnetic tape drives or a track of a magnetic tape" of claims 12 and 16, in the Abstract, Stolowitz discloses a disk drive array. Further, in Figure 5, Stolowitz discloses a multiplexer (510) for changing the data from a parallel state to a serial state. However, Stolowitz doesn't explicitly disclose a magnetic tape having data blocks and a parity block in which the data blocks and the parity block are serially arranged on the magnetic tape with the parity block following the data blocks and the parity block being based on the data blocks. On pages 176-177, White discloses that the format of a QIC tape typically contains 20 to 32 parallel tracks. Each track is divided into blocks of 512 or 1,024 bytes, and segments typically contain 32 blocks. Of the blocks in a segment, eight contain error-correction codes. These tracks comprise both data and parity. It would have been obvious to one of ordinary skill at the time of the invention to use a magnetic tape in the redundant storage system of Stolowitz. A person of ordinary skill in the art would have been motivated to make the modification because in column 8, lines 32-33, Stolowitz

discloses the use of a serial stream when reading from the disk drives which is necessary for a tape drive. Thus there is a means for changing parallel data from the disks to serial data. The serial data from a tape would also be able to be inputted into the system of Stolowitz. Further, on pages 12-13 and in Figure 5, the Applicant discloses that an obvious variation of the magnetic tape drive is an array of disks. Specifically, on page 13, lines 5-7, the Applicant discloses that *controller 18 writes to and reads from storage elements in the same manner as described with reference to the track of magnetic tape 14 in FIG. 3.* Also, the system of Stolowitz is compatible with a tape disk drive because it contains a SCSI bus, which is a common interface for devices such as CD-ROM, drives and backup tape drives as well as hard disks (see column 4, lines 15-17).

Referring to the remaining limitations in claims 12 and 16:

a. On page 176, White discloses that of the blocks in a segment, eight contain error correction codes (the parity block following the data block).

b. In the Abstract, Stolowitz discloses a disk drive array with parity data based upon data blocks and a disk drive array controller that carries out disk drive data transfers.

c. In column 6, lines 20-22, Stolowitz discloses methods and circuitry for effecting synchronous data transfer to and from an array of disk drives (reading blocks sequentially from respective data storage elements).

d. In column 8, lines 42-44, Stolowitz discloses reconstructing missing data in the event of any single drive failure. Further, in column 8, lines 32-37,

Stolowitz discloses that missing data is reconstructed as the serial stream of read data moves from the drives into the buffer. Only, complete, correct data is stored into the buffer according to the invention. No delay is incurred in the process. Hence a bad sector (corrupted or unreadable) or even an entire bad drive causes no special read delay (determining if the data block currently being read is good or bad based on the reading of the data block currently being read).

e. In column 6, lines 37-63, Stolowitz discloses to execute a read operation, a single global read command comprising the starting sector number, the number of sectors to be read, and a "read sector" command type is "broadcast" to all of the drives at once. After all of the devices are ready, i.e., the requested data is ready in each drive buffer, that data is transferred from all of the drive ports to the buffer memory using a single sequence of common strobes which are shared by all of the drives. This data in turn is read out of the buffer to the host (providing the data block currently being read to the host if the currently being read data block does not follow a bad data block).

f. In column 8, lines 50-55, Stolowitz discloses that once the data from the last drive enters the pipeline, the accumulator will be holding the data from the missing drive. This result is transferred to a hold latch, and when the missing word in the pipeline from the failed drive is reached, the contents of the hold latch is substituted in place of the pipeline contents (if one of the data blocks is bad, storing the good data blocks following the bad block in sequential order).

g. In column 8, lines 42-48, Stolowitz discloses that to reconstruct missing data in the event of any single drive failure, the serialized read data stream is passed through an N+1 stage pipeline register. To begin, a word from the first drive is loaded into an accumulator and into the pipeline. As the next data word enters the pipeline from the next drive, it is XORed with the first word and the result stored in an accumulator.

h. In column 8, lines 50-55, Stolowitz discloses that once the data from the last drive enters the pipeline, the accumulator will be holding the data from the missing drive (reading the parity block from the magnetic tape after all of the data blocks have been read).

i. In column 8, lines 48-55, Stolowitz discloses that the accumulating process is repeated for each subsequent drive except that data from the failed drive is ignored. Once the data from the last (redundant) drive enters the pipeline (reading the parity block from the parity storage element), the accumulator will be holding the data from the missing drive. This result is transferred to a hold latch, and when the missing word in the pipeline from the failed drive is reached, the contents of the hold latch is substituted in place of the pipeline contents (if one of the data blocks is bad, reconstructing the bad data block from the accumulated parity of the good data blocks and the parity block in order to form a reconstructed good data block; providing the reconstructed good data block to the host; and providing the stored good data blocks to the host in sequential order after the reconstructed good data block has been provided to the host).

j. In column 8, lines 42-44, Stolowitz discloses an N+1 stage pipeline register (a buffer for storing the good data blocks read by the controller after the bad data block until the controller reconstructs the bad data block to preserve ordering of the data blocks during reading).

Referring to claim 13, in column 8, lines 42-48, Stolowitz discloses that to reconstruct missing data in the event of any single drive failure, the serialized read data stream is passed through an N+1 stage pipeline register. To begin, a word from the first drive is loaded into an accumulator and into the pipeline. As the next data word enters the pipeline from the next drive, it is XORed with the first word and the result stored in an accumulator (accumulating parity of the good data blocks includes exclusive ORing the parity of the good data blocks read prior to the good data block currently being read with the good data block currently being read).

Referring to claim 14, in column 8, lines 48-55, Stolowitz discloses that the accumulating process is repeated for each subsequent drive except that data from the failed drive is ignored. Once the data from the last (redundant) drive enters the pipeline, the accumulator will be holding the data from the missing drive. This result is transferred to a hold latch, and when the missing word in the pipeline from the failed drive is reached, the contents of the hold latch is substituted in place of the pipeline contents (reconstructing a bad data block includes exclusive ORing the accumulated parity of the good data blocks and the parity block).

Referring to claim 17, in column 8, lines 42-48, Stolowitz discloses that to reconstruct missing data in the event of any single drive failure, the serialized read data

stream is passed through an N+1 stage pipeline register. To begin, a word from the first drive is loaded into an accumulator and into the pipeline. As the next data word enters the pipeline from the next drive, it is XORed with the first word and the result stored in an accumulator (the parity accumulator accumulates parity of the good data blocks by exclusive ORing the parity of the good data blocks read prior to the good data block currently being read with the good data block currently being read).

Referring to claim 18, in column 8, lines 48-55, Stolowitz discloses that the accumulating process is repeated for each subsequent drive except that data from the failed drive is ignored. Once the data from the last (redundant) drive enters the pipeline, the accumulator will be holding the data from the missing drive. This result is transferred to a hold latch, and when the missing word in the pipeline from the failed drive is reached, the contents of the hold latch is substituted in place of the pipeline contents (reconstructing a bad data block includes exclusive ORing the accumulated parity of the good data blocks and the parity block).

(11) Response to Argument

On page 8 of the Appeal Brief, under section 3. The Claimed Invention Compared to Stolowitz and White, the Applicant argues, "Accordingly, White does not teach or suggest data blocks and a parity block serially arranged on a track of the magnetic tape with the parity block following the data blocks in the manner of the RAID-4 standard as claimed. In contrast, White teaches blocks serially arranged on a magnetic tape track with some of the blocks being data blocks and some of the blocks being data and ECC blocks." The Examiner respectfully disagrees. At the bottom of

page 176 in an illustration, White that the EC blocks are grouped together and that they are not both data and EC blocks. Further, on page 176 in point 5, White discloses that each track is divided into blocks of 512 or 1,024 bytes, and segments typically contain 32 blocks. Of the blocks in a segment, eight contain error-correction codes. At the top of page 177 in an illustration, White shows DATA bits followed by EC bits. Further, on page 177 in point 2, White teaches a buffer used to prepare data that is written onto the tape and discloses that if the tape drive's controller includes chips that handle error correction, the backup software dumps the full buffer from RAM to the controller's own buffer, where the chips append error correction (EC) codes. If the controller doesn't have built-in error correction, the software computes the EC codes based on the pattern of 0 and 1 bits in the files, *appends them to the end of the data* (emphasis by Examiner) in the RAM buffer, and copies the contents of the RAM buffer to the controller buffer. For these reasons, the Examiner submits that White does teach data blocks and a parity block serially arranged on a track of the magnetic tape with the parity block following the data blocks in the manner of the RAID-4 standard as claimed, and it is believed that the rejections should be sustained.

Respectfully submitted,

Robert Beausoliel

MCM
October 7, 2004

ROBERT BEAUSOLIEL
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100

TIMOTHY R SCHULTE
STORAGE TECHNOLOGY CORPORATION
ONE STORAGETEK DRIVE MS 4309
LOUISVILLE, CO 80028-4309

DONALD SPARKS
SUPERVISORY PATENT EXAMINER